Structure-Significant Representation of Structured Datasets

Raghu Machiraju, *Member*, *IEEE*, Zhifan Zhu, *Member*, *IEEE*,
Bryan Fry, and Robert Moorhead, *Senior Member*, *IEEE*

Abstract—Numerical simulation of physical phenomena is now an accepted way of scientific inquiry. However, the field is still evolving, with a profusion of new solution and grid-generation techniques being continuously proposed. Concurrent and retrospective visualization are being used to validate the results, compare them among themselves and with experimental data, and browse through large scientific databases. There exists a need for representation schemes which allow access of structures in an increasing order of smoothness (or decreasing order of significance). We describe our methods on datasets obtained from curvilinear grids. Our target application required visualization of a computational simulation performed on a very remote supercomputer. Since no grid adaptation was performed, it was not deemed necessary to simplify or compress the grid. In essence, we treat the solution as if it were in the computational domain. Inherent to the identification of significant structures is determining the location of the scale coherent structures and assigning saliency values to them [22], [23]. Scale coherent structures are obtained as a result of combining the coefficients of a wavelet transform across scales. The result of this operation is a correlation mask that delineates regions containing significant structures. A spatial subdivision (e.g., octree) is used to delineate regions of interest. The mask values in these subdivided regions are used as a measure of information content. Later, another wavelet transform is conducted within each subdivided region and the coefficients are sorted based on a perceptual function with bandpass characteristics. This allows for ranking of structures based on the order of significance, giving rise to an adaptive and embedded representation scheme. We demonstrate our methods on two datasets from computational field simulations. Essentially, we show how our methods allow the ranked access of significant structures. We also compare our adaptive representation scheme with a fixed blocksize scheme.

Index Terms—CR Categories and Subject Descriptors: I.3.2 [Computer Graphics]: Graphics Systems; I.3.8 [Computer Graphics]: Applications; I.4.2 [Image Processing]: Compression (Coding), wavelet transform, structure detection, human visual system, progressive transmission.

1 Introduction

Given the advances in hardware and the proliferation of new numerical and data analysis techniques, the solution grids for computational simulations are continuously becoming larger. For instance, even 100 time steps using a grid size of 69 × 43 × 15 from a computational field simulation (CFS) requires the storage of the computational grid, a vector field (velocity), and two scalar fields (density, energy), all in floating point format, making the storage requirement 136 Mbytes. Both concurrent and retrospective visualization are employed to understand the results of a simulation. The size of the datasets is a burden, especially when they are transmitted over a slow network. It is not uncommon that the domain for a simulation is spatially divided into several zones and executed on a multicomputer or a cluster of networked processors. Each zone has its own geometric grid and a designated processor, or solver. Visualizing the solution data from the solvers in real-time or near-real-time provides instantaneous feedback and, therefore, allows tracking of the simulation. Such systems are becoming increasingly available [9], [20]. In this case, visualization is done in a distributed manner, i.e., the solution data are transferred from the solvers to a graphics workstation over the network and selectively visualized. Therefore, it is useful for the datasets to be stored and transmitted in a compressed format. Even more important for solutions which have not converged yet is a quick meaningful preview so that the simulation can be steered.

A coarse-to-fine progressive display is often used. However, it may be useful to display the most significant structures first. In Fig. 1, we show an isosurface rendering of pressure around a turbine. The structures are ranked in terms of singularity; it is useful to see structure A in its entirety before structure D, which is far smoother. It is easier to gauge the shape of D than of A and, hence, a coarser rendition of D is sufficient. The structures are displayed again in Fig. 1b using a rate (percentage of complete information) of 20 percent. The structure A is faithfully rendered while structure D is quite distorted. The *embedded visualization* paradigm allows a ranking of structures and can be very useful for concurrent visualization. This paradigm of visualization is different from displaying a progressively smoother version of the complete rendering. We believe that progressive refinement is less useful for scientific visualization. Given the loss of detail

- *R. Machiraju, B. Fry, and R. Moorhead are with the Department of Computer Science, Mississippi State University, Mississippi State, MS 39762 and with the MSU NSF Engineering Research Center for Computational Field Simulation. E-mail: {raghu, bfry, rjm}@erc.msstate.edu.*
- *Z. Zhu is with Duke Engineering and Services, 1180 Town Center Dr., Las Vegas, NV 89134. E-mail: Zhifan_Zhu@notes.ymp.gov.*

For information on obtaining reprints of this article, please send e-mail to: tvcg@computer.org, and reference IEEECS Log Number 106601.

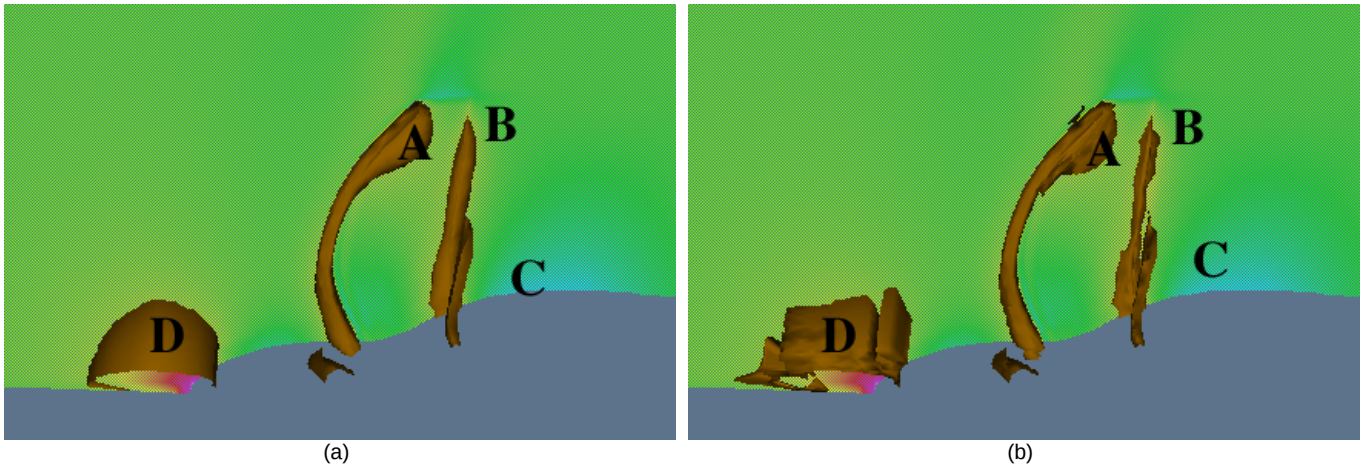

(a) (b)

Fig. 1. Structure significant visualization. (a) Isosurface rendering of pressure around a turbine is depicted. The structures A, B, C, and D are ordered in terms of increasing smoothness. (b) Isosurfaces rendered with only 20 percent of wavelet coefficients.

in the subsampled images, features cannot be easily discerned. A computational scientist cannot learn much about the physical phenomenon from a subsampled image. Also, it is not clear that the bandwidth is efficiently utilized.

The needs for retrospective visualization are equally stringent. Given the increased use of computational methods to solve engineering and scientific problems, large archival collections of datasets now exist and terascale visualization is no longer a concept but a need for scientists working on large-scale problems. There is a need, therefore, to index, browse, and retrieve through a scientific database. The indexing can be in the form of a coarse-to-fine structures, while browsing and retrieval require the display of structures as described in Fig. 1. Also, computational scientists often need to compare the results of a simulation with another (Fig. 2; same physical condition with different solvers) and with experimental data. Thus, methods which allow side-by-side comparison of datasets based on the structures will be useful. Even more useful will be a method which will display progressively refined structures based on some parameter-like rate. Once again, an embedded scheme can be employed to enormous benefit. Transform coding schemes have been used for efficient transmission and compressed storage of images. The discrete cosine transform and the wavelet transform have been used for this purpose. Both lossy and lossless schemes have been used to store and transmit datasets. The new paradigm of visualization proposed here can also be supported by the same transform coding and multiresolution framework.

We now examine some characteristics of CFS datasets which require special consideration when designing representation schemes:

- CFS datasets are generally very sparse, since the computational experiment attempts to study and focus on very specific phenomenon and, sometimes, at very specific locations in the domain. Most of the information is packed in a very small portion of the entire volume, irrespective of the grid topology used.
- CFS datasets sometimes have sharp singularities, like a shock. However, as Fig. 2 shows, the singularities are not very sharp, and a method like edge detection will fail to detect regions of activity. Wavelet methods, on the other hand, are often able to locate the significant structures.
- For computational field simulations, it is important that any representation be inherently lossless. However, it is also imperative that the compression scheme allow for selective display of important structures in a dataset.
- Finally, CFS datasets come in a variety of shapes and sizes, often dictated by the underlying geometry and the desired resolution of structures. Many multiresolutional schemes impose *power-of-two* restrictions on the resolution. A commonly used solution is padding the dataset until *power-of-two* resolutions are obtained. This increases the quantity of data that must be transmitted, sometimes significantly. Thus, it may be worthwhile to partition the dataset into blocks wherein the wavelet transform can be applied.

1.1 Proposed Method

To satisfy all our requirements of concurrent and retrospective visualization, we make use of a technique which identifies regions of significant structures in datasets. The structures we locate are scale coherent, i.e., contain all frequencies across scales. Our method is two-pass in nature. In the first pass, a spatial subdivision is obtained which delineates regions of high spatial frequency. The subdivision is obtained by creating a correlation mask from the wavelet transform. This step may require temporary padding of the volume to create power-of-two dimensions. An octree embedding of this mask provides the required subdivision. In the second phase, each block of the partition is coded separately using an appropriate function from a wavelet library. This step does not require any padding of the original volume. Later, the wavelet coefficients in each of the blocks are sorted and packed based on the presence of significant structures. The visual content associated with coefficients within each block is estimated by employing a computational model of the Human Visual System (*HVS*). The final result is an embedded encoding of the volume wherein a

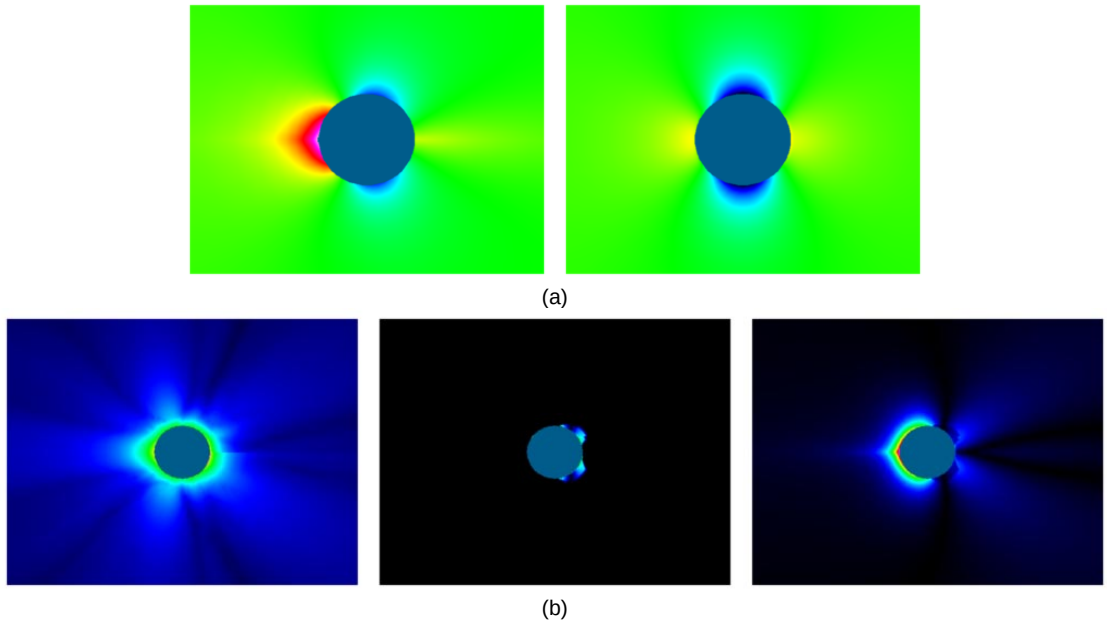

Fig. 2. (a) Density around a right circular cylinder computing using first (left) and second order (right) upwind methods. (b) Density computed around an airfoil at 0 (left), 1 (middle), and 4 (right) degree angles of attack. (Courtesy: John West, DoD Major Shared Resource Center, U.S. Army Corps of Engineers, Vicksburg, Mississippi.)

ranking of coefficients is possible. Our multiresolution representation scheme can be very easily extended to lossy or lossless coding. Similarly, it allows for:

- browsing of datasets allowing selective and embedded display,
- comparative visualization of datasets,
- lossy / lossless compression of floating point data,
- progressive transmission, allowing the user to stop transmission at any desired time.

Our attention to the problem arose when some scientists were executing computational programs on a supercomputer at a very remote site and viewing the results on workstations. The computational problem was the study of hydrodynamics around submarine hulls (dataset 2). A structured curvilinear grid was employed to model the domain around the submarine. Given the very slow link (17 kilobytes per second) and the availability of a supercomputer, it seemed imperative to employ a representation scheme that was amenable to efficient bandwidth utilization and provided useful visualization. Useful implementation on a slow link is achieved when the most important information is transmitted first. The grid was sent separately or was already known a priori at each end of the link. It was only the actual solution that had to be compressed for this particular situation. In essence, we treat the solution as if it were in the computational domain. Certainly, this method will not apply to unstructured meshes. This approach of not including the grid geometry is far from ideal, especially if grid simplification is the desired goal. On the other hand, simplification entails the elimination of certain grid cells and will require resampling. In computational simulations, unless the grid is adapted to the solution at various epochs, it is perhaps moot or even incorrect to simplify the grid. Convergence of the solution is often a very difficult task and uncontrolled simplification can destroy the integrity of the solution.

In Section 2, we motivate the need for a variable block-sized multiresolution representation and review past reported methods in volume compression. In Section 3, we present our scheme for multiresolution volume representation and, in Section 4, we present our embedded progressive transmission scheme, which allows a perceptual ranking of structures. Results are presented in Section 5 and, in Section 6, we draw conclusions and discuss future work.

2 Background

In this section, we describe the coding and analysis capabilities of the wavelet transform. Later, we describe a particular computational model of the human visual system. Our proposed method includes this model to perceptually rank structures in a dataset. The words signal, image, and volume are used synonymously in this paper. Also, the words compression and coding are employed to describe a representational technique.

Previous volume compression methods include predictive methods [11], fractal methods [3], vector quantization (VQ) [28], Discrete Cosine Transform (DCT) [42], wavelets [9], [13], [16], [26], [27], [32], [34], [38], [39],[40], and Gaussian pyramids [12]. The predictive method (followed by Huffman coding) reported in [11] provided sufficient improvement over the standard UNIX utilities like gzip, but did not provide a high level of compression. Also, it was not amenable to embedded coding for progressive transmission. Ning and Hesselink [28] used VQ to compress volume data and achieved a high level of compression. However, the method is again not very amenable to concurrent progressive transmission given the asymmetric nature of the coder and decoder. The coding effort is generally orders of magnitude higher than expended in the decoding component. The fractal technique of [3] is expensive and is not amenable to progressive transmission. Among the

transform coding methods, the DCT, wavelet, Gaussian, and Laplacian image pyramids have found common use. However, the work reported in both [12] and [42] emphasizes less the compression aspect and more the rendering of images from the compressed domain. There exist other methods which have been borrowed from image coding and are suitable for hardware implementation [19].

We now describe multiresolution representation schemes based on wavelets. Our focus is on schemes which are

1) adaptive,
2) embedded, and
3) assign perceptual saliency and, hence, ranking to structures when employed in conjunction with a simple computational model of the HVS.

Such schemes will allow one to achieve the goals of efficient concurrent and retrospective visualization. Adaptive methods which are adept at capturing higher-order smooth (beyond edges) singularities and can locally adjust in terms of spatial resolutions can provide more efficient representation schemes.

2.1 First-Generation Wavelet Representations

In transform coding (TC) methods, the image or volume is projected to a transformed space spanned by adequate basis functions. The goal of all TC methods is to decorrelate the signal and scale the energy in the projected space so that a few coefficients can adequately represent the signal. The wavelet transform has gained much popularity due to some very attractive properties, including compact support of the basis functions and the ability to represent functions of arbitrary smoothness [6]. Analysis and compression can be obtained by retaining only a portion of the coefficients. The size of the retained portion is dictated by the desired quality of the resulting image or the smoothness of the function. It has been shown in [8] that an image can be suitably represented by the first N coefficients within a certain error tolerance measured in spaces spanned by smooth functions. Methods which retain only N coefficients are often classified as *first-generation methods* [9], [13], [26], [32], [34], [38]; they do not attempt to characterize the behavior from the subband information. Rather, they apply standard and computationally viable techniques for compression and transmission.

2.2 Second-Generation Wavelet Representations

As pointed out in [10], second-generation methods are driven by structures in an image. Moreover, it is essential that a human observer or a model of the HVS be included in the coding scheme. These methods attempt to locate information in an image and then proceed with coding the coefficients which exist in the identified areas. Much information in an image is concentrated in a few regions. These are dominated by structures characterized by discontinuities in intensity and manifest as edges (images) and boundaries (volumes). Mallat and Zhong [25] used spline wavelets to obtain a redundant wavelet decomposition. The end result of their nonorthogonal transform is a hierarchy of local absolute maximas and, hence, multiscale edges. These maximas are not searched for; instead, a numerical procedure is used to characterize the Lipschitz coefficients of singularities. Westermann [39] implemented the same ideas to detect structures in a time-varying 3D volume. Edge-based methods suffer from the disadvantage of being error prone in the presence of noise or when the image has a profusion of small features or texture. A significant method, wavelet probing, was presented in [5] to segment the image into smooth regions which are coded very efficiently. The identified edge regions are then coded separately by employing the same technique. Although the method may yield high coding efficiency, it does not facilitate the rankings of structures as described in Fig. 1. The proposed method fits in this same category.

2.3 Adaptive Methods

Adaptive methods which are adept at capturing higher-order smooth (beyond edges) singularities and can locally adjust, in terms of spatial resolutions, can provide more efficient representation schemes. A compression and progressive transmission scheme for visualization should maximize the information presented while minimizing the data rate. This objective can be explained by using an information-rate curve. If the information in a data volume is uniformly distributed over the data domain and the bit budget is uniformly allocated for each point, this curve will be linear. In most cases, however, the information contained in a dataset is nonuniformly distributed, resulting in a nonlinear curve. The desired information-rate curve should have a sharp rise as close to the zero-rate as possible, so that most of the data information can be visualized as early as possible. This is a departure from the rate-distortion perspective normally espoused in the image coding literature. To achieve the requirements of concurrent and retrospective visualization, an embedded scheme can be used, where the information is ranked by perceptual significance and transmitted accordingly over the channel [32], [43]. In [30], a zerotree wavelet structure is created from the wavelet space. This allows for a grouping of wavelet coefficients of similar significance in a hierarchical fashion and facilitates an embedded representation that is well suited for progressive transmission. In [18], this technique was expanded to include a computational model of the HVS. Said and Perlman extended this algorithm and obtained superior results [29].

Since the wavelet transform is performed using a single mother wavelet, the transform is essentially signal independent. To adapt to the statistical properties of the data, wavelet packet techniques [4] decompose the signal by choosing from a library of wavelet packets the wavelets that best represent the signal. However, the computational expense is still staggering and the wavelet packet transform is far from robust. An aspect of the wavelet packet transform can be usefully employed; spatial subdivision of the image is conducted and, normally, a quad-/octree is used in conjunction with the subdivision of the frequency bands. Each block of the spatial subdivision can be coded with a different wavelet basis, the choice being dictated by a variety of reasons including measures of function smoothness.

Thus, we have described the basic rationale for our method and outlined the methodology. Two essential elements are missing, namely a technique to detect structures

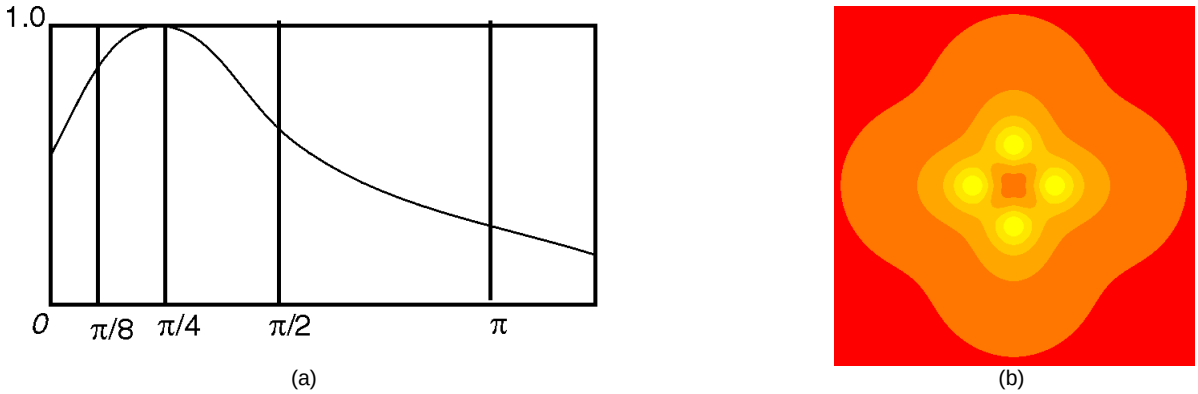

(a) (b)

Fig. 3. (a) One-dimensional contrast sensitivity function. The dyadic subdivision (of the wavelet transform) is imposed on the image frequency. (b) Two-dimensional version of the *CSF*.

and a description of a computational model of the HVS. We describe a methodology in Section 3 for detecting structures across scales and end this section with a description of the Contrast Sensitivity Function (CSF) and a computational model of the HVS.

2.4 Human Visual System (HVS)

The human cortex is often modeled as a linear system and its response to a visual excitation in the receptive field is expressed as a convolution of the impulse response of the visual cortex with the input stimulus. Equivalently, the response to an input can be obtained by multiplying the Contrast Sensitivity Function (CSF) with the Fourier transform of the input. The CSF measures the response of the visual system to different frequencies. It is important to note that the response is lower for higher frequencies. Attempts have been made to measure the CSF of the human visual system. For instance, Mannos and Sakrison [21], after conducting a series of psychophysical experiments on human subjects, found that the CSF can be modeled by the function in (1) (Fig. 3a). Here, f_s is spatial frequency in cycles per degree.

$$C_s(f) = \left[0.0492 + 0.2964 f_s\right] e^{-\left(0.114 f_s\right)^{1.1}}. \quad (1)$$

A 2D (or a 3D) version can be easily obtained and is shown in Fig. 3b. The discrete orthogonal wavelet transform does not serve as a computational model for the HVS. To include the CSF into our methods, a frequency decomposition induced by a dyadic wavelet transform is imposed on the CSF (Fig. 3a). Thus, for each band, a weight can be computed which is then applied to the wavelet transform. The coefficients at each level and location are modulated using a perceptual weight obtained by integrating the CSF over the area occupied by the band in a 3D frequency space. Fig. 3a shows the CSF function mapped to image frequency and the superimposed wavelet transform subdivision in 1D. Thus, after m levels of the wavelet transform, the weight is given as:

$$C_m = \frac{\left(\int_{FB_m} C_s(\omega) d\omega\right)}{A\left(FB_m\right)}, \quad (2)$$

where FB_m is the frequency subband $\left(\frac{\pi}{2^m}, \frac{\pi}{2^{m-1}}\right)$ in question and $A(FB_m)$ is the width of the band. The weights are precomputed and stored in a table. In Fig. 7b, we report the normalized weights for a 2D subdivision [18]. The weights are based on the assumption that the image is viewed at a distance six times the image height. In the following section, we show how structures can be adaptively detected.

3 Adaptive Structure Detection

The volume representation scheme reported herein has the four components shown in Fig. 4. The block representation allows for the ranking and measurement of important information in a signal. The progressive transmission scheme requires another component to order the blocks (Section 4). The first three components facilitate the detection of significant structures, while the fourth component achieves the actual coding. We now describe each component in detail.

3.0.1 Wavelet Transforms

We used both the dyadic wavelet transform and the nonorthogonal wavelet transform of Mallat and Zhong [25]. The biorthogonal wavelet transform is employed to code the image in each block. The choice of the nonorthogonal transform allows for a better characterization of the local spatial frequencies of the underlying function. Moreover, it does not suffer from the aliasing problems that the critically sampled dyadic transform suffers from. The result of the wavelet transform is a pyramid of detail subvolumes and a smooth subvolume. The subband volumes are of the same size as the original volume.

In [35], a study of different biorthogonal wavelets for coding efficiency was conducted. Our choice of the second wavelet was dictated by that study. We chose the biorthogonal filter of second order since it provides adequate results. The reported method should work for any volume at any resolution. However, its effectiveness depends on the partitioning of the volume into regions which are then coded individually. Padding is required for resolutions which are not powers-of-two for the first wavelet transform. Once the partition is obtained, the padded regions are discarded and only the original volume is considered. More details on the wavelet transform can be found in much reported work including [6], [24]. It now remains to identify the coefficients which contribute to structures in an image.

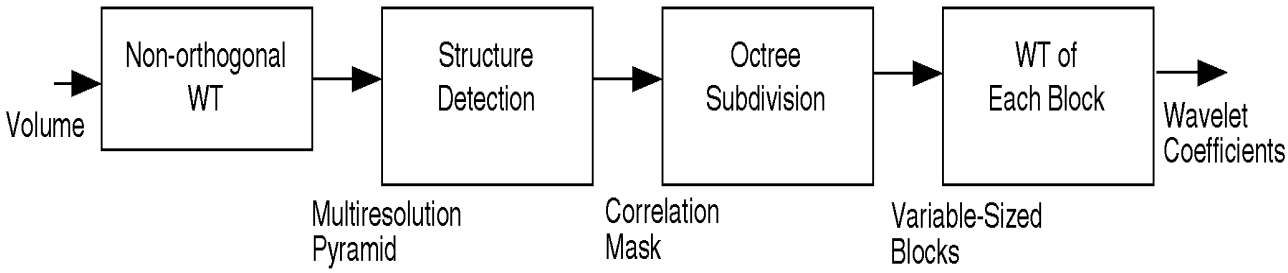

Fig. 4. Schematic diagram of coding algorithm. The volume is subjected to a wavelet transform (WT). A correlation mask is found which is then subdivided by an octree. The ensuing blocks are then coded.

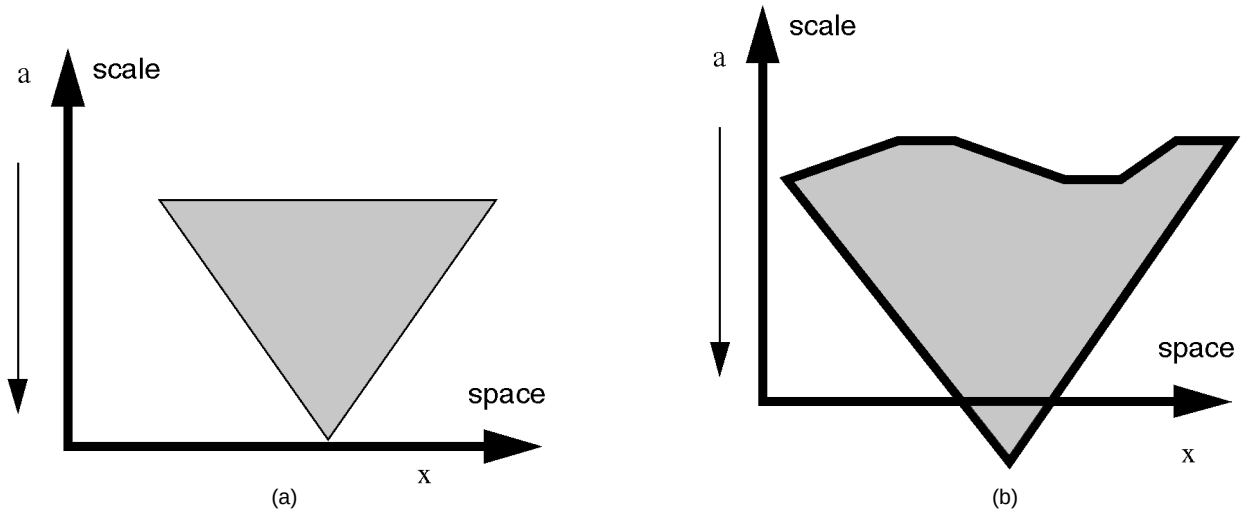

Fig. 5. Cone-of-influence. The *y*-axis measures scale in a decreasing fashion. (a) *Delta Dirac* Function—The half-plane is a cone, not always symmetrical. The shape is really dictated by the wavelet function. (b) Sums of Many Deltas—The cone is shifted below the *x*-axis.

We now provide some motivation of our basic structure detection method. More details can be found in [22], [23].

3.1 Wavelets as an Analysis Tool

The wavelet transform provides valuable information for the *space-frequency* diagram or the *scalogram* [6], [15]. A space-frequency diagram describes the frequency content of the image at a pixel location at different scales of resolution. The impulse and the infinitely long sinusoidal wave are extreme examples. The space-frequency diagram for the impulse signal is defined at only one instance in the space domain (localized), while it completely occupies the frequency range (not localized). The situation is changed when the sinusoidal signal is considered. Only one location of the frequency spectrum is occupied (localization), while the signal persists over all points in space. We now informally describe the shape of the space-frequency diagram of any signal. We derive it from the shape of a corresponding diagram of an impulse.

3.1.1 Space-Frequency Diagram of an Arbitrary Signal

The wavelet transform of a *Delta Dirac* function δ centered at $x = 0$ with respect to a wavelet function ψ is given as:

$$W[\delta : \psi](a,b) = \frac{1}{a}\psi\left(\frac{b}{a}\right). \tag{3}$$

The space frequency diagram (*a*–scale parameter; *b*–time/space parameter) thus has all frequencies. The set of points in the half plane forms the so-called *cone-of-influence* (Fig. 5a). Now, consider a function composed of the sum of n deltas located at $x = x_n$, each scaled differently by a factor k_n. From the property of linearity [15] of the wavelet transform, we get

$$W[f_n : \psi](a,b) = \sum_n \frac{k_n}{a}\psi\left(\frac{b - x_n}{a}\right). \tag{4}$$

The cone-of-influence is now shown in Fig. 5b. The cone is shifted toward the smaller scales such that the bottom is outside the half plane. Since all signals can be expressed as the summation of infinitely many Dirac functions, it can be said that the wavelet transform of all signals has a shifted cone form. It is our goal to determine the intersection of the cone with the x-axis. The significant structures populate only this spatial region. The rest of this section describes the process of locating them and assigning saliency to them.

3.1.2 Combining Function

The combining function C is defined as follows [22], [23]: For a J-level wavelet decomposition W, at all x, C exists and is given by

$$C(x, J, W) = \left|A_J(x)\right| \prod_{j=1}^{J} \left|D_j(x)\right|, \tag{5}$$

where A_j and D_j, $j = 1 \ldots J$, are the subband functions at location x. Thus, we obtain a floating-point mask which has the

same size as the image. We now state some observations about the mask and discuss scale-coherent structures.

3.1.3 Scale-Coherent Structures

$C(x, J, W) = 0$ implies the absence of structures at that location. This arises from the cone-like shape of the wavelet transform. Where $C(x, J, W) > 0$, coherent structures exist and lie within the support of the cone. The effectiveness of the method presented in this paper rests on the following observations:

- As J becomes large and approaches infinity (if possible), the result of the combining function should approach the support of the continuous wavelet transform. The idea behind the combining function is simple. The smallest scales have the smallest supports and, hence, by multiplying across scales one obtains a size equal to the size of the smallest support. Similar observations and arguments were made for images elsewhere [1].
- It is possible that a nondyadic powers-of-two transform will yield a better mask, since the space-frequency diagram is measured at a large number of scale positions. Thus, a more accurate mask can be obtained if an M-band (M > 2) spectral subdivision is conducted.
- In Fig. 6, we show the effect of combining. The masks in Fig. 6b and Fig. 6c have been greatly exaggerated for display purposes. The former is obtained through the use of the orthogonal transform, while the latter is obtained from the application of the nonorthogonal transform of [25]. The mask in Fig. 6c is smoother. The masks are from two different yet nearby locations which contain similar structures. The product is high at locations near the point of discontinuities. The product locates correlated structures across scales. The number of scales, the size of the structures, and the wavelet functions used also contribute to the effectiveness of the combining algorithm.
- Noise reduction occurs in the product. Uncorrelated additive noise in the spatial domain becomes correlated in the wavelet domain [15]. However, the size of the correlation is smaller in the coarser scales. Thus, in the smooth space measured by Aj, uncorrelated noise is suppressed and the coefficients arising from the noise process measure close to zero.
- If f is Lipschitz (α), the function f is defined as

$$\left|f(x) - P_n(x)\right| \le Kx^{\alpha} \qquad n \le \alpha \le n+1.$$

Essentially, the function f differs from an n-degree polynomial by a bounded oscillatory fraction. All known practical functions are known to satisfy this definition of a Lipschitz function. The wavelet coefficients at level j are bounded [17] as shown:

$$\left|W[f](j,x)\right| \le K2^{j\alpha}. \tag{6}$$

The combining function is then bounded by

$$C(x, J, W) \le A^{S(J)}2^{\alpha S(J)}, \quad S(J) = \sum_{1}^{J} j, \tag{7}$$

where the function $S(J)$ is the sum of the positive integers up to J. Other functions may be used instead of the product. The logarithm may also be employed. In that case, the combining function is bounded as

$$\log\left(C(x, J, W)\right) \le S(J)\log A + \alpha S(J). \tag{8}$$

Since the mask depends on the value of α, the value of the mask is an indication of the significance of the singularity. Larger values of α indicate presence of significant singularities. Hence, the above arguments for the effectiveness of the mask in capturing singularities are valid.

3.2 Scale-Coherent Structure Detection

Edge detection captures all the zero crossings of the second derivative of a function; a mask delineates all high-frequency regions. If the wavelet had a regularity of *one*, then the combining mask would contain mostly edges. Our method employs simple strategies to combine information across scales and is rooted in the approximation theory perspective of the wavelet transform [15]. Starck et al. [31] report a combining method which employs the sum of binary thresholded wavelet coefficients to create a mask. The binarization process loses important detail and can introduce blockiness. We, however, explore the use of the product function. The product of the wavelet transform is a more natural choice, since it arises from the shape of the wavelet transform. The location of structures and the assignment of saliency values is, therefore, more accurate. Xu et al. independently developed an algorithm based on the product [41]. Their motivation was based on the fact that correlation of structures is best captured by the product. Also, they attempt to find edges, rather than regions of arbitrary singularity. Our motivation stems from the detection of structures, and denoising is a component in the process. By using the product of the wavelet transform, we obtain a mask which identifies the location of structures and assigns them saliencies. The wavelet transform of the image and its inverse transform form the first and last stages of the algorithm, respectively. The other components of the algorithm extract a minimal possible representation of the image in terms of coherent structures.

3.2.1 Denoising

CFS datasets are not tainted with as much noise as those from medical scanners. Uncertainty arising from computational inaccuracies, number representations (truncation, overflow, and underflow), and grid tolerances does exist. The ensuing noise can be suitably modeled with a Gaussian distribution function. This characterization, although not paramount to the functioning of many numerical schemes, does allow for better location of features. As described in Section 4, the process of combining inherently denoises. It is, however, useful to denoise even before generating the subband mask. That way we have a minimum number of coefficients that go into the construction of the mask. We resort to the statistical thresholding methods of Donoho and Johnstone [7]. We implement an adaptive thresholding scheme, which determines the threshold at different levels from a noise model.

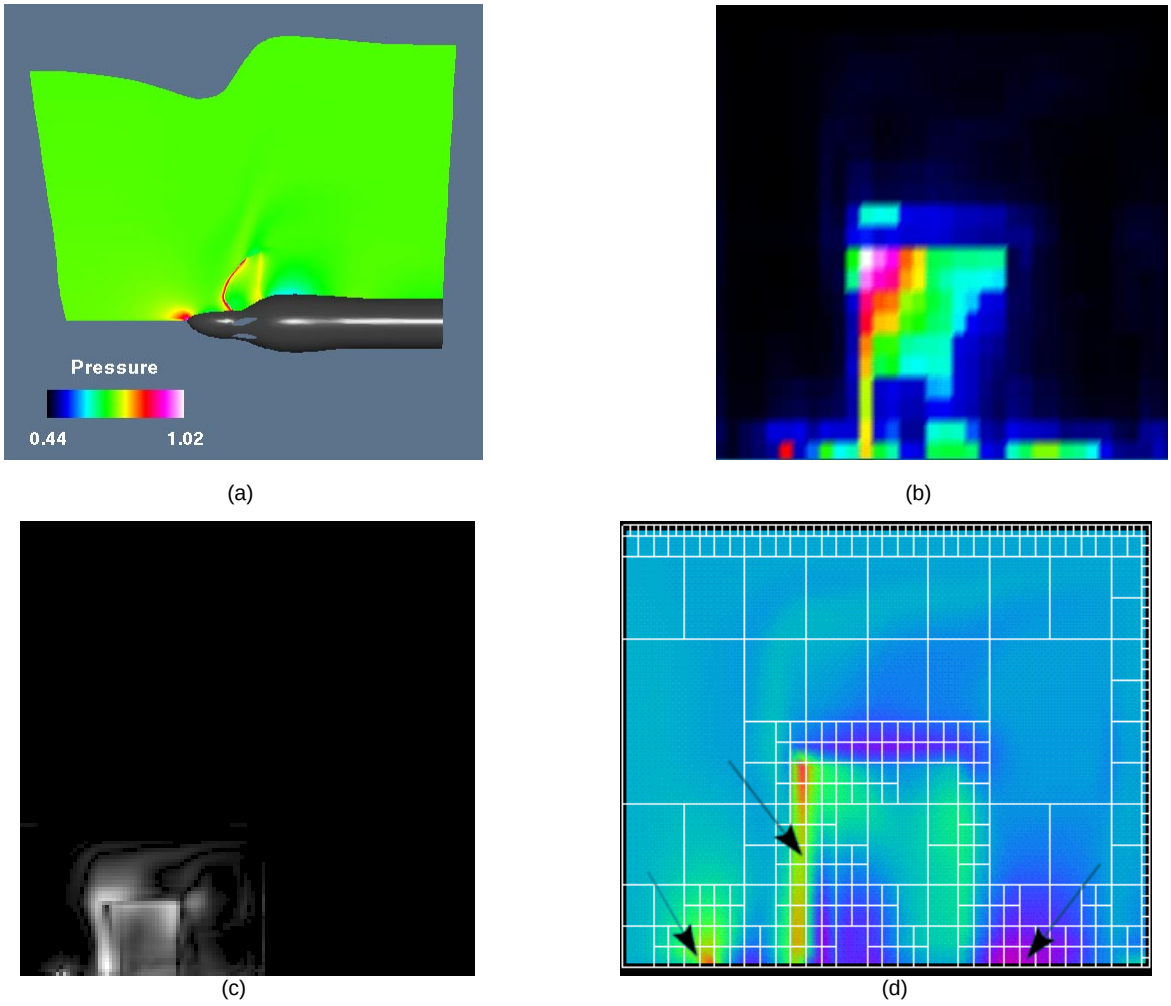

Fig. 6. Detection of scale coherent structures and their enhancement. (a) Simulation of flow around a turbine blade (source: Mark Janus, ERC, MSU). (b) Correlation (color quantized using the spectral color table) mask for a 2D slice in computational domain using the orthogonal transform. (c) Correlation (8-bit grayscale) mask using the nonorthogonal transform of a slice near the first one. The masks are mostly sparse, with large values near regions of activity. The orthogonal mask is more blocky. (d) Spatial subdivision imposed on the original.

3.2.2 Feature-Based Thresholding

More coefficients can be eliminated based on the relative magnitude of the product mask. A typical discriminating function, shown in (9), is employed to remove insignificant features and to shrink the dynamic range. We describe the equation for images. The quantity $C(i, j)$ is the mask value at pixel (i, j), C_{max} is the maximum value of the mask, g is a tunable parameter, and R is the current dynamic range. The value of g controls the number of structures (and, hence, coefficients). Values between 1.1 and 1.5 have been found to be acceptable for most images. Higher values of g degrade the visual quality of the image.

$$d = R \times \log\left(1 + \left(\frac{C(i,j)}{C_{max}}\right)^{g}\right) / \log(2). \qquad (9)$$

Wherever the subband mask is set to zero, the corresponding coefficients in the pyramid have to be set to zero. In other words, the subband mask has to be propagated back into the multiresolution pyramid. As a result of this thresholding operation, only the coefficients contributing to coherent features remain. Since the subband images are the same size as the original image, the thresholding becomes trivial.

3.3 Octree Subdivision of the Mask

The spatial subdivision allows the identification of coherent regions in the image. Several options exist, including region growing and recursive spatial decomposition methods. Region growing algorithms are not very amenable to user control and can result in blocks of odd shape and size. Hence, we resort to recursive spatial subdivision methods based on quad-/octrees. The level of the tree is controlled by the user, thus producing square/cubic blocks. Otherwise, a crude homogeneity measure can be employed to control the subdivision process. A threshold measuring the amount of information is used. The subdivision process stops when the amount of information is small as measured by the user threshold. Thus, a larger threshold would yield larger blocks. We show the spatial subdivision obtained using a slice of the data in Fig. 6d. There will be some blocks which do not carry any information; they arise in the proximity of the discontinuities. The effectiveness of the adaptive partitioning scheme is made evident in Section 5 when the performance is compared against fixed block partitioning schemes. In the following section, we describe how the wavelet coefficients in the blocks can be packed for progressive transmission, i.e., our embedded representation scheme.

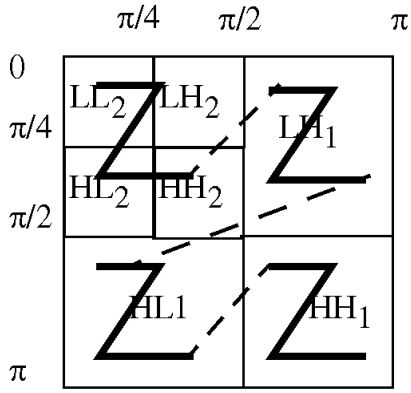

(a)

0 \ π/4 \ π/2 \ π	π/4	π/2	π
0	0.1363	0.1310	0.0516	0.0516
π/4	0.1310	0.1017	0.0516	0.0516
π/2	0.0516	0.0516	0.0219	0.0219
π	0.0516	0.0516	0.0219	0.0219

(b)

Fig. 7. Embedding of dyadic subbands. (a) The subband LL_1 has been split into four more subbands. Each of the four subbands of the two-level transform has an associated coefficient. The selection of the bands and within bands is shown in the form of Z. Thus, the order among the subbands is LL_1, LH_1, HL_1, and HH_1. In lieu of LL_1, the subbands LL_2, LH_2, HL_2, and HH_2 are embedded. (b) The normalized weights for two-level decomposition of the *CSF*. The weights are used in the computation of the remaining information (RH_i) function.

4 Ranking of Structures

In this section, we describe the packing of data for progressive transmission. The blocks obtained from the spatial subdivision of the mask are sorted based on the information content within each block. A single bitstream is created from the wavelet transform coefficients of the partitioned blocks for transmission or storage. The rationale behind this scheme is to insert the most vital information at the beginning of the bitstream. Such a scheme then allows for transmission until the desired target rate is achieved. The information content can be measured in a variety of ways, including the mean-square error. However, the MSE is not a viable measure of saliency in an image [33]. Rather, a weighted MSE based on perceptual considerations is gaining much popularity and is being used in image coding extensively [37]. The resulting bitstream is constructed from coefficients obtained from blocks in an order determined by the presence of significant structures.

We employ a 3D version of the CSF to code a given volume. In reality, there is no physical justification for such a function. However, it is not uncommon to employ modulation transfer functions (MTF) to alter the significance of coefficients in a transformed space. Thus, we can conceivably construct a device which accentuates frequencies in the central portion of a 3D spectrum. To avoid unnecessary confusion, we, however, use the terms CSF and MTF synonymously. The perceptual weighting of wavelet coefficients is commonly done in the image coding literature. The modulation with a perceptual function allows one to give some bands higher weights. Thus, if an inverse wavelet transform is performed and a visualization is performed, the structures which are perceptually significant are enhanced, while the less significant ones are subdued. The pertinent visualization technique can influence perception of the phenomenon. However, a choice and ranking of structures based on perceptual significance does not impair the perceptual abilities of the viewer.

In reality, we need a function which allows us to rank the coefficients. It seemed convenient to use the perceptual function. Our weighting scheme is an approximate one. The perception modulation function should be applied to the images; however, we apply it early on in the visualization pipeline. In the case of parallel projection, the application of the 3D function is not improper. In reality, the 3D function is a simple extension of a 2D function; a parallel projection will only alter a scaling factor but not the essential characteristics of this function. We now show how the perceptual weight can be incorporated into a packing scheme.

4.1 Perceptual Weighted Embedded Packing Scheme

The following scheme is designed to select the next significant coefficient based on the structure information estimated in each block. As a result, a wavelet coefficient bitstream is constructed with the coefficients carrying the most information being placed in the stream as early as possible. After data partitioning, the data volume is represented by N variable-size blocks, B_i, $0 \le i \le N$. Each block B_i has the following entities:

- The starting spatial address (i, j, k) and the block size.
- The information quantity contained in the block, H_i, which is defined as the sum of all correlation mask values contained in the block.
- The wavelet coefficients computed for the data block.

The measurement of the information in a block does not have to be exact. As long as H_i is a monotonic function, it will suffice. Since the mask is created from the product and nonlinearly enhanced with a logarithm function, H_i will be monotonic. Blocks of the same size will have the same CSF weight but different H_is. Thus, the visual significance of a wavelet coefficient is a relative quantity and is determined by two factors:

- the quantity of information in the block to which the wavelet coefficient belongs, namely H_i,
- the position of the wavelet coefficient inside the block.

During embedding, the relative significance of the blocks and coefficients has to be constantly updated. To do so, we calculate for each block RH_i, the remaining information, using a CSF which is precomputed for each possible block size. This function should have the following properties:

$$RH_i(N_i) = 1$$
$$RH_i(0) = 0, \tag{10}$$

where N_i is the number of coefficients in a block.

Thus, a $2k \times 2k \times 2k$ block will have $8k^3$ coefficients, each

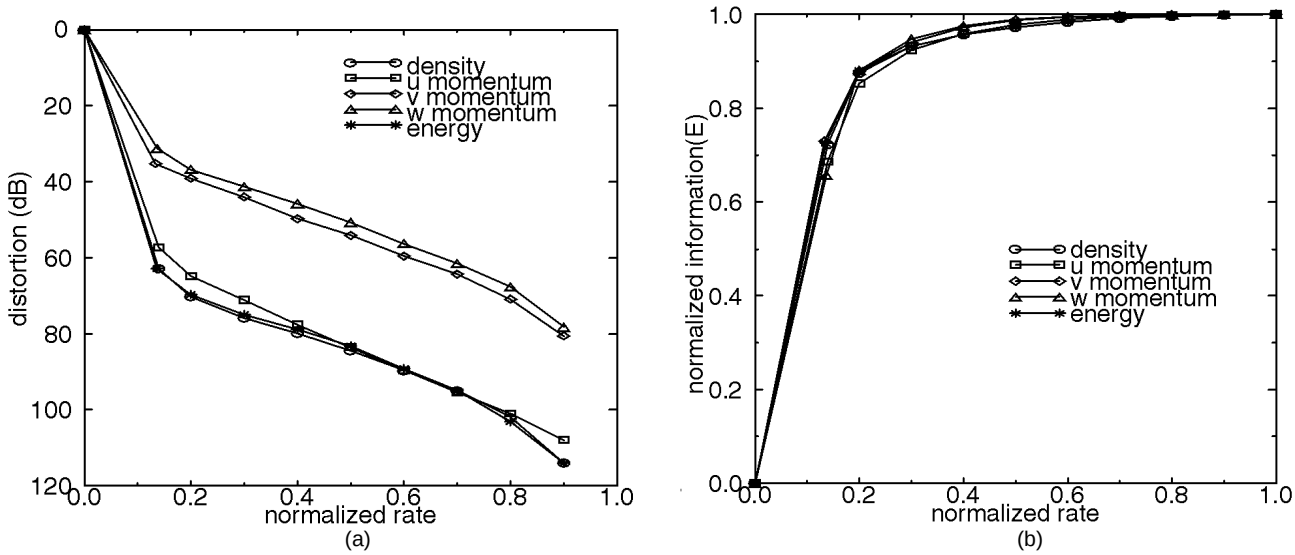

Fig. 8. Performance of our coding scheme. (a) Rate-Distortion curve. The distortion is measured in decibels and is given by $-20 log_{10}$(Error). It is plotted along the Y-axis and the rate along the X-axis. (b) Rate-Information curve. The information in each block is measured by the sum of the mask values.

measuring the information available in a subband. In the case of a two-level, 2D wavelet transform, the dyadic decomposition results in the subdivision shown in Fig. 7. Each band in the dyadic decomposition is precomputed and assigned a value, i.e., each position in the block (i, j, k) is assigned a normalized weight $NCSF(i, j, k)$ (prefix N indicating normalization). A possible candidate for the function RH is

$$RH_i(m) = H_i\left(1 - \sum NCSF(i, j, k)\right) \quad m \in Embedded. \tag{11}$$

The function RH essentially computes the perceptual information that can still be assigned to the block after m coefficients have been embedded. When a coefficient is embedded in the bitstream the corresponding normalized weight is subtracted from the value of RH. The coefficients inside a block are selected in a zig-zag order similar to the selection scheme in JPEG [36]. Once a coefficient is embedded, the corresponding perceptual weight is subtracted. This scheme distinguishes between blocks that have the same spatial dimension but different information content. Given our octree spatial subdivision of the original domain, several such similar sized blocks can exist, some with little information. The packing algorithm can be summarized as follows:

Step 1: Sort the blocks in descending order, using H_i as the key.

Step 2: Insert into the bitstream the lowest subband coefficient for each block and update RH_i accordingly. Resort the blocks based on RH_i.

Step 3: Insert into the bitstream coefficients (determined by the Z patterns) from the first block in the sequence and update RH_i until the first block is no longer the most significant block.

Step 4: Insert-sort the first block into the block sequence based on RH_i; go back to Step 3. In [29], [30], a variant of this embedded packing scheme is described.

The block addresses and the ordering information must be coded and stored precisely first in the bitstream. This can be done either by: ·

- coding the block addresses in the ordered sequence, or ·
- creating a "bit-volume" and putting the block's sequence number at the block's reference address in the volume and, then coding the bit-volume.

We now provide the results of applying our coding scheme to a dataset from a simulation.

5 Results

We tested our methods on two datasets. The first one is obtained from a simulation of the nonsteady flow of fluid between turbine blades. The flow of fluid is axisymmetric. Fig. 6a shows a slice of the density field juxtaposed with the body of the turbine [16]. The regions of interest are the leading and trailing edges of the blade and the tip of the propeller. The grid is curvilinear and is of size 69 × 43 × 15. The solution of the simulation yielded two scalar fields (density, ρ; energy, e) and a vector field (u, v, w). We do not compress the grid in this exercise, rather we restrict ourselves to the scalar fields and the components of the velocity field. Thus, the grid file and the compressed solution comprise the data that needs to be transmitted or stored. To determine the correlation mask, we create a larger grid of size 128 × 64 × 16 through zero-padding. An octree subdivision of the mask provides the required blocking. Starting with the original data in each block, we then perform a wavelet transform. Finally, the coefficients and the blocks are all embedded into a bitstream. The orthogonal wavelet transform was employed to obtain the embedded representation.

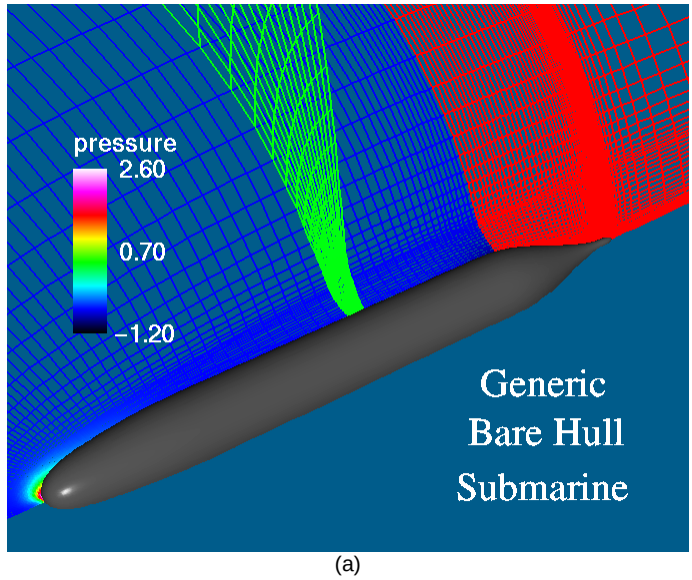

(a)

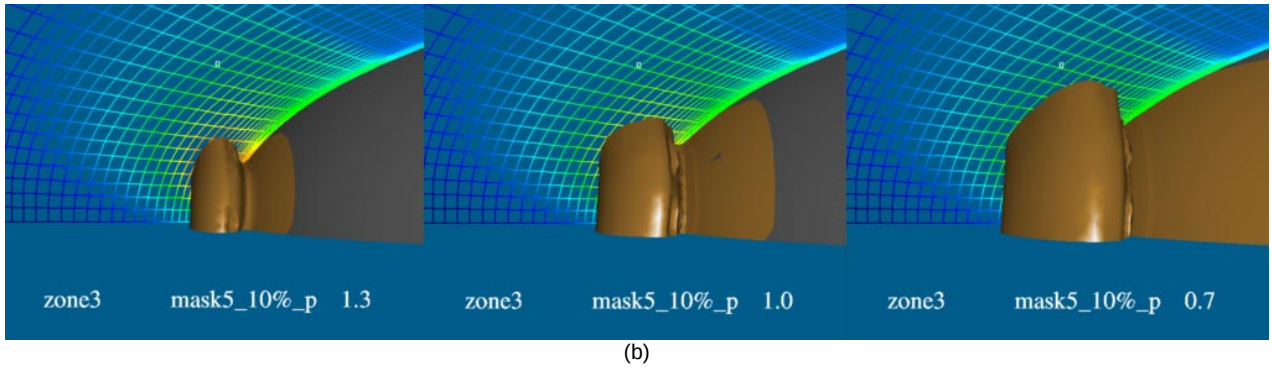

(b)

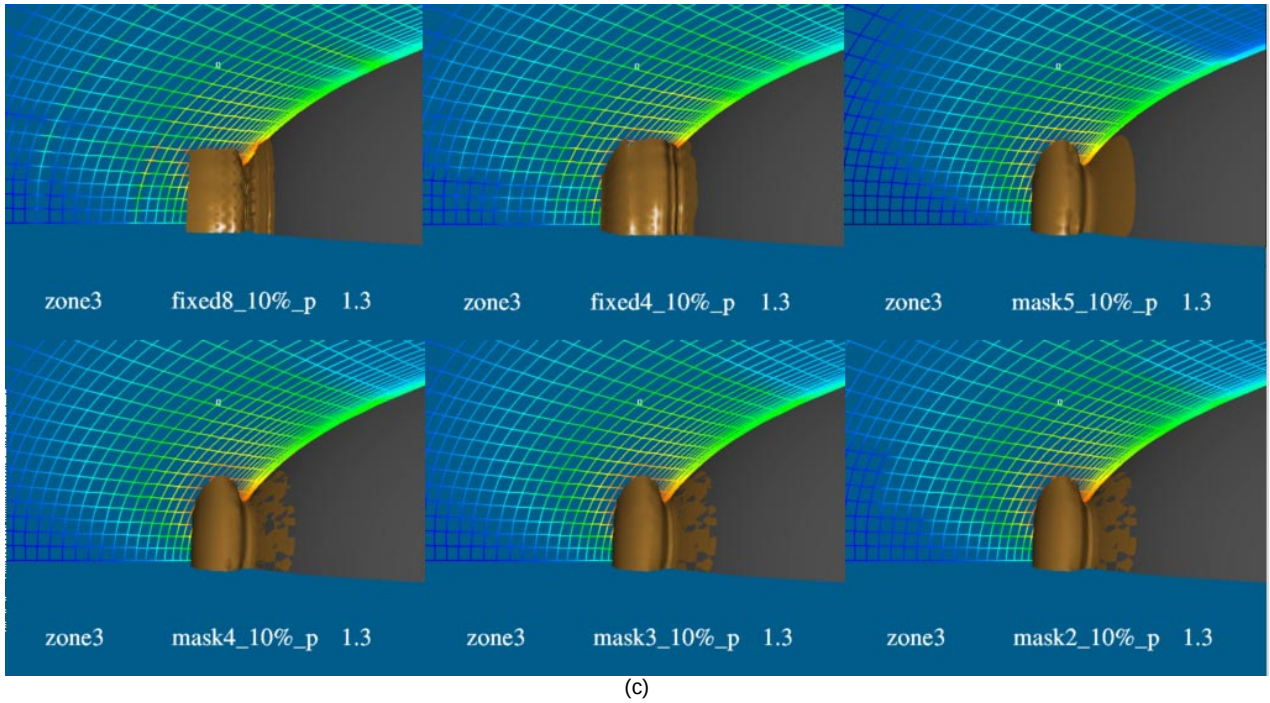

(c)

Fig. 9. Flow around a submarine hull. (a) Basic configuration. The complete geometry has eight zones, four in front and four behind. The pressure variations are shown on slice planes for various zones of the curvilinear grid. (b) For the same octree subdivision of the mask (threshold = 0.5) and rate (= 10%), pressure thresholds of 1.3, 1.0, and 0.7 are visualized in zone 3 (front of submarine). (c) The efficacy of the adaptive partition is shown. The isosurface for the pressure value of 1.3 and rate 10 percent is displayed for a variety of partitioning strategies. Cubes of size 4^3 and 8^3 are used and compared with partitions obtained when thresholds of 0.5, 0.4, 0.3, and 0.2 are used. It is clear that adaptive partitions are superior, especially for larger threshold values. Fixed partitions result in blocky isosurfaces.

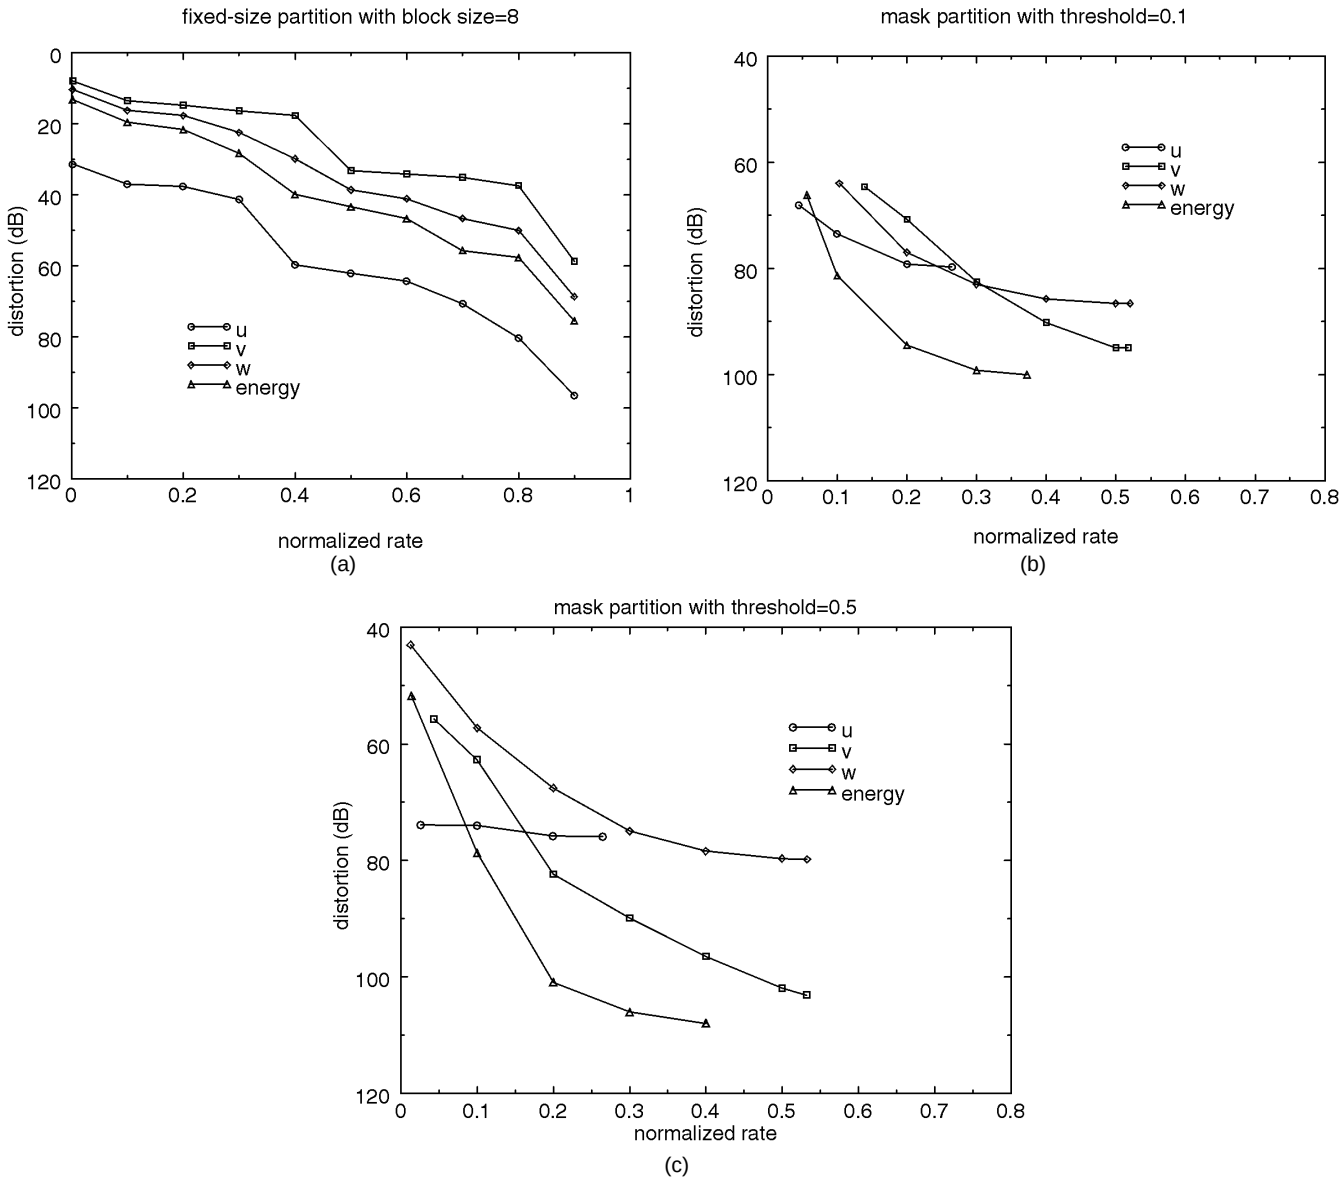

Fig. 10. Rate distortion curves for dataset 2. (a) Fixed mask partition of size 8^3. Adaptive partition with (b) low threshold (c) high threshold.

We selected another dataset (Fig. 9a) on which to test our technique. Since it was of mutual interest, we chose to attack concurrent visualization of an evolving flow. The volume is divided into 16 zones and horizontal symmetry is assumed so that the flow is only computed in eight zones, four in the front and four in the back, each zone spanning a 45 degree cylindrical wedge. Each zone is of size 81 × 73 × 17. The purpose of the simulation was to determine integrated quantities, such as thrust and torque, at various angles of attack. In the picture, a cutting plane in the k direction (k = 1 for zone 1) is shown with the pressure (a derived quantity) mapped onto it. A i slicing plane at i = 60 in zone 1 is shown in green. A slicing plane at k = 1 in zone 2 is shown in red. The colormap used in all of the zone 3, timestep 2 (*qq32*) data is also shown. We now discuss the rate-distortion and rate-information curves for the test datasets. The nonorthogonal wavelet transform was employed to determine the correlation mask for the second dataset.

5.1 Rate-Distortion Curve

The rate-distortion curve is plotted for the components of the first dataset in Fig. 8a. We define the rate to be the percent of the coefficients used. Thus a rate of one would imply that all coefficients are used, while a rate of zero implies that none are used. We plot the distortion for all the fields. The distortion at a given rate is measured by finding the difference between the actual and the partially reconstructed dataset in the mean-square-error (MSE) sense. It is evident that the degradation is nonuniform over the various fields. The v and the w-component of the velocity yield similar behavior; they require that 80 percent of the data be used to obtain a very desirable error of 60dB. Given the axisymmetric flow, most disturbances are not in the principal flow direction. In reality, this is a very encouraging statistic. The other fields fare even better, and the same distortion is obtained at the even lower rate of 20 percent. Essentially, the distribution of structures is very sparse, and the datasets are amenable to extreme compression.

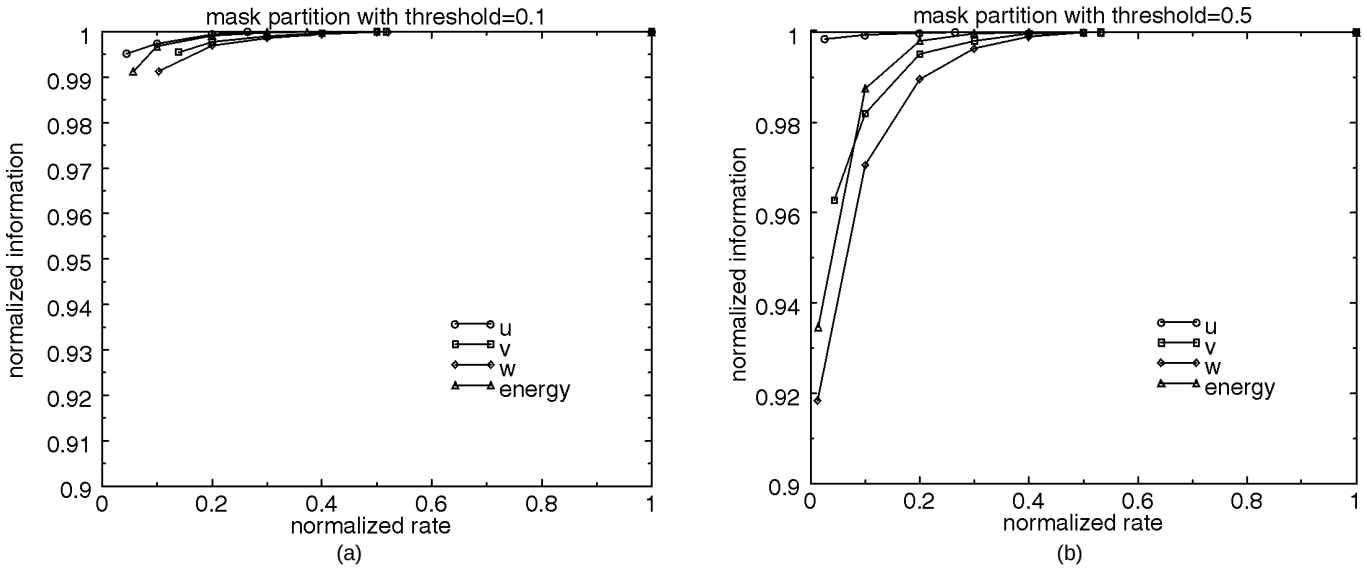

Fig. 11. Information-curve-performance of adaptive partitioning. (a) Threshold = 0.1, (b) Threshold = 0.5.

We also compared the rate-distortion performance for different partitioning schemes. The results for dataset 2 are plotted in Fig. 10. One can see that the distortion among the various quantities is more spread than in dataset 1. A reason for this observation could be the nonconvergence of the solution. The adaptive schemes show more graceful degradation. The partition with higher threshold (0.5) has more desirable characteristics. The parameter u (the horizontal velocity) exhibits seemingly anomalous behavior, since the distortion is constant for the range of values plotted. A plausible explanation is that the submarine's motion is along the same direction and, thus, this component of the velocity is constant. The images in Fig. 9c corroborate the rate-distortion curves.

5.2 Information-Rate Curve

We also plot the information that is embedded in the bitstream at different rates. The information embedded is given by the quantity E, defined as:

$$E = 1 - \frac{\sum_i RH_i}{\sum_i H_i}. \tag{12}$$

Both the quantities H_i and RH_i are defined in Section 4. The denominator in the ratio measures the information content and is computed as the sum of the mask values over all the blocks. In essence, it measures the information in terms of significant structures. The numerator, on the other hand, measures the information retained.

Although the rate-distortion curve is different for at least two of the components, the information-rate curve is similar for all of the parameters of dataset 1 (Fig. 8b). This is the case since we measure information in terms of structures as captured by the mask. The masks of all five components are very similar, signifying that the regions of discontinuities are similar. Since the quantity E is a gross measure of perceptually significant structures in a dataset, the quantity evaluates to the same value for all components. However, the number of significant blocks is much less for the v, w components, indicating that significant structures in these datasets are much sharper or discontinuous. Thus, the rate-distortion curve is less steep and requires more coefficients to reconstruct. For dataset 2, we show the rate-information curve for the adaptive schemes (Fig. 11). The lower threshold seems to facilitate a better characterization of the information content.

5.3 Progressive Isosurface Rendering

To show the efficiency of our approach, we display the isosurfaces for a derived quantity, namely pressure. We use FAST [2], a tool often used by CFS engineers to visualize their solution datasets. (Note: The checkerboard pattern in the background of many of the images is due to a well-known artifact of the screen door transparency rendering process used in FAST.) Note that we did not take into account the uncertainty introduced by the visualization process itself. For incompressible flow, pressure at any point can be computed as follows:

$$p = (\gamma - 1)\left[e - \frac{1}{2}\rho\left(u^2 + v^2 + w^2\right)\right], \tag{13}$$

where γ is a gas constant.

5.4 Dataset 1

For a pressure of 0.75 (the range of pressure is [0.44, 1.02]), Fig. 1a shows the isosurfaces at the full normalized rate. We see a profusion of surfaces of varying smoothness and curvature. The isosurfaces at a rate of 20 percent are shown in Fig. 1b. It is again worthwhile noting that the highly curved surface in the middle is captured with high fidelity, even at 20 percent. The smoother surface at the extreme right is captured accurately at the rate of 30 percent (not shown). It is important to note that, even at a rate of 20 percent, the isosurface along the blade is captured at very high fidelity. This is a testimony to the fact that our packing scheme was effective in embedding significant structures very early in the bitstream. Another important observation is that the behavior of our scheme was not affected, even though we

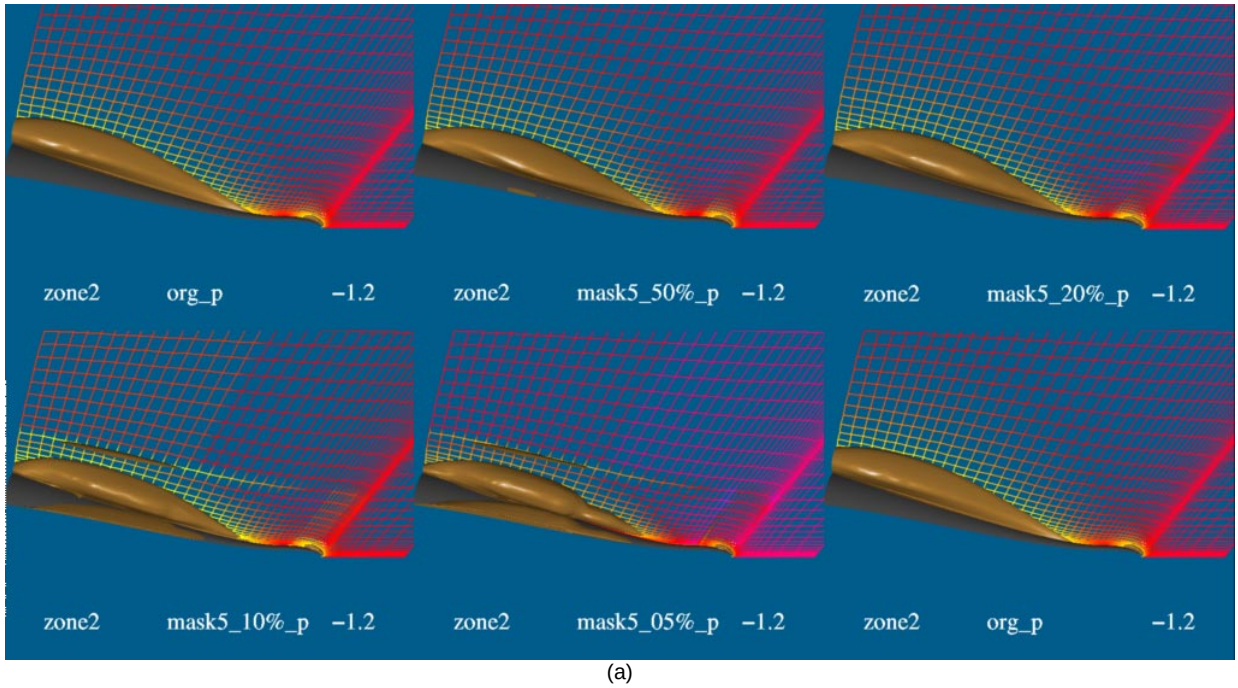

(a)

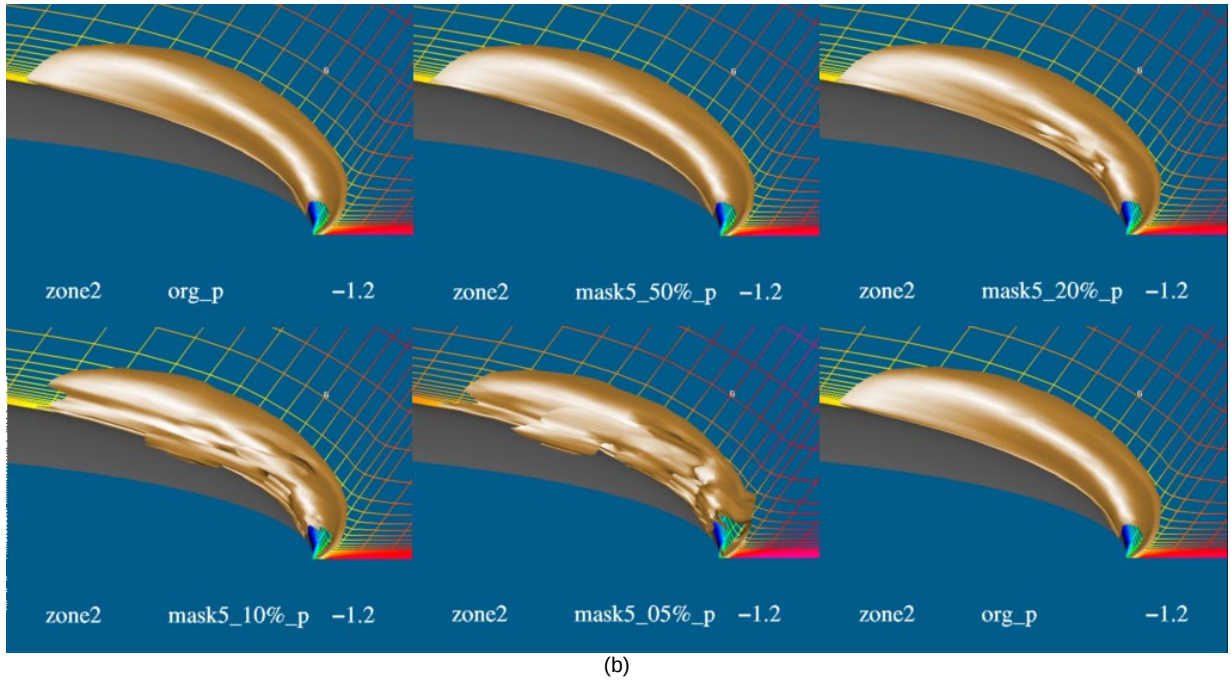

(b)

Fig. 12. Isosurface renderings of pressure (value = –1.2) of zone 2. (a) Different rates and same adaptive partition. (b) The lower images are zoomed versions of the above.

applied our methods on a nonlinear derived quantity, namely pressure. The results contained in [9] do not show the same amount of structure-sensitivity; instead, the degradation at rates less than full is more uniform. In [9], CFS datasets are compressed using a nonadaptive blocking scheme.

5.5 Dataset 2

Since the solution has not converged, there is no one correct pressure value to visualize. Thus, we need to make some reasonable choice. In Fig. 9b, we show the isosurfaces for pressure values of 1.3, 1.0, and 0.7. The spatial extent of the isosurface is inversely proportional to the pressure. The pressure ranges from –1.2 to 2.6 within the volume. The small white circle is a glyph at $i = 13$, $j = 33$ to provide some indication of the size of the isosurface plume in grid coordinates, since we do not want the shape of the isosurface to be influenced by the interpolating errors. We will use the 1.3 value for further presentation herein. The results we report are consistent for all three isosurface values, but only the results for pressure of 1.3 are shown for clarity.

The question naturally arises as to how much improvement there is with the masking and variable size blocks. In Fig. 9c, we demonstrate just how much. Sending only 10

percent of the coefficients, the isosurface for pressure of value 1.3 produces a rendition very close to that obtained when using all of the coefficients. Note that both fixed size block levels (4^3 and 8^3) produce poor reproductions of the correct plume. In zones 2, 4, 6, and 8, there exist two low pressure regions along the body. This is shown in Fig. 12 for zone 2. Note the isosurface quality, even with only 10 percent of the coefficients retained. The smaller isosurface at the very end of the body is visible in the lower right of each panel in Fig. 12a, near the dense nearly-vertical grid lines. In Fig. 12b, a zoomed version of the isosurface in the low pressure region formed at the trailing end of the submarine is shown. The glyph is at $i = 43$, $j = 23$, so the isosurface slices approximately 20 × 15 grid cells. The colors are ordered as for the pressure field in zone 3, but the range in zone 2 is from –2.64 to -0.17.

The encoder we used is designed for testing and the timings for it are not indicative of how fast it could be. The nonoptimized decoder runs in about 10 seconds on a Silicon Graphics Indigo 2 with a 100 Mhz R4400 CPU. Compressing floating point numbers to achieve high compression rates is not a trivial task. Usually, an entropy scheme, like Huffman coding or even vector quantization, can be employed to compress the wavelet coefficients [9]. It suffices to say that any further compression and quantization schemes in the wavelet domain will only enhance the efficiency of the reported scheme. The decoding cost is also small given the regular pattern of embedding the coefficients and the fact that biorthogonal wavelets were chosen to actually encode the blocks.

6 Conclusions and Future Work

We presented a multiresolution representation of a volume that is amenable to high compression ratios and progressive transmission. We implemented a better method of correlation mask generation than that used in our earlier work [43]. Essentially, we replaced the dyadic orthogonal transform with a nonorthogonal one. Results on two significantly different datasets are shown. Finally, we compared the performance of the adaptive partitioning methods with fixed block methods. The results obtained bear testimony to the value of the method. The method is computationally cheap and can include aspects of the Human Visual System easily. Improvements can come in the form of

- better methods to determine the correlation mask,
- better packing schemes,
- adaptive choice of the wavelet function from a library.

New work can include:

- further quantization of the coefficients, ·
- detecting coherency between successive frames in a time varying simulation,
- datasets from other domains (medical) and modalities,
- compression of geometry.

Acknowledgements

We acknowledge the contributions of Prof. Roni Yagel, The Ohio State University and Dr. Ajeet Gaddipatti of GE Medical Systems toward the development of the wavelet combining schemes. One of us (Raghu Machiraju) developed the schemes in collaboration with them. This work was supported by the MSU National Science Foundation Engineering Research Center for Computational Field Simulation, Mississippi State University. Special thanks also go to Prof. Mark Janus, Dr. Ramesh Pankajakshan, and Dr. Shyam Neerarambam of the ERC for access to several datasets.

References

[1] J.-P. Antoine, P. Carrette, R. Murenzi, and B. Piette, "Image Analysis with Two-Dimensional Continuous Wavelet Transform," *Signal Processing*, vol. 31, pp. 241-272, 1993.

[2] G. Bancroft, F. Merritt, T. Plessel, P. Kelaita, R. McCabe, and A. Globus, "FAST: A Multi-Processed Environment for Visualization of Computational Fluid Dynamics," AIAA Paper 91-0793, *Proc. 29th Aerospace Sciences Meeting*, Reno, Nev., Jan. 1991.

[3] W.O. Cochran, J.C. Hart, and P.J. Flynn, "Fractal Volume Compression," *IEEE Trans. Visualization and Computer Graphics*, vol. 2, no. 3, pp. 313-322, Dec. 1996.

[4] R.R. Coifman and M.V. Wickerhauser, "Entropy-Based Algorithms for Best-Basis Selection," *IEEE Trans. Information Theory*, vol. 38, no. 2, pp. 713-718, Mar. 1992.

[5] B. Deng, B. Jawreth, G. Peters, and W. Sweldens, "Wavelet Probing for Compression Based Segmentation," *Wavelet Applications in Signal and Image Processing*, A.F. Laine, ed., pp. 266-276, *Proc. SPIE 2034*, 1993.

[6] I. Daubechies, *Ten Lectures on Wavelets, CBMS-NSF Regional Conf. Series Applied Mathematics*. Philadelphia: SIAM, 1992.

[7] D.L. Donoho and I.M. Johnstone, "Ideal Spatial Adaptation via Wavelet Shrinkage," *Biometrika*, vol. 81, pp. 422-455, 1993.

[8] R.A. DeVore, B. Jawreth, and B.J. Lucier, "Image Compression Through Wavelet Transform Coding," *IEEE Trans. Information Theory*, vol. 28, no. 2, pp. 719-746, Mar. 1992.

[9] X.S. Du and R.J. Moorhead, "Multiresolutional Visualization of Evolving Distributed Simulations Using Wavelets and MPI," *Proc. SPIE*, vol. 3,017, *Proc. SPIE/IST Electronic Imaging '97*, pp. 298-309, San Jose, Calif., Feb. 1997.

[10] J. Froment and S. Mallat, "Second Generation Compact Image Coding with Wavelets," *Wavelets—A Tutorial in Theory and Applications*, C.K. Chui, ed., pp. 655-678. Academic Press, 1992.

[11] J. Fowler and R. Yagel, "Loseless Compression of Volume Data," *Proc. 1994 Symp. Volume Visualization*, pp. 43-50, Washington D.C., Oct. 1994.

[12] M.H. Ghavamnia and X.D. Yang, "Direct Rendering of Laplacian Pyramid Compressed Volume Data," *Proc. Visualization '95*, pp. 192-199, Oct. 1995.

[13] M. Gross, L. Lippert, A. Dreger, and R. Koch, "A New Method to Approximate the Volume Rendering Equation Using Wavelets and Piecewise Polynomials," *Computers Graphics*, vol. 19, no. 1, pp. 47-62, 1995.

[14] B. Guo, "A Multiscale Model for Structure-Based Volume Rendering," *IEEE Trans. Visualization and Computer Graphics*, vol. 1, no. 4, pp. 291-301, Dec. 1995.

[15] M. Holschneider, *Wavelets: An Analysis Tool*. Oxford Univ. Press, 1995.

[16] J.M. Janus, H.Z. Horstman, and D.L. Whitfield, "Unsteady Flowfield Simulation of Ducted Prop-Fan Configurations," AIAA-92-0521, *Proc. 30th Aerospace Sciences Meeting and Exhibit*, Reno, Nev., Jan. 1992.

[17] S. Jaffard, "Wavelets and Nonlinear Analysis," *Wavelets: Mathematics and Applications*, J.J. Benedetto and M.W. Frazier, eds., pp. 467-503. Boca Raton, Fla.: CRC Press.

[18] Y. Kim, I. Cho, I. Lee, T. Yun, and K.T. Park, "Wavelet Transform Image Compression Using Human Visual Characteristics and Tree Structure with a Height Attribute," *Optical Eng.*, vol. 35, no. 1, pp. 204-212, Jan. 1996.

[19] G. Knittel, "High Speed Volume Rendering Using Redundant Block Compression," *Proc. Visualization '95*, pp. 176-183, Oct. 1995.

[20] L. Lippert, M.H. Gross, and C. Kurmann, "Compression Domain Volume Rendering for Distributed Environments," *Proc. Eurographics '97, Computer Graphics Forum*, vol. 14, no. 3, 1997.

[21] J.L. Mannos and D.J. Sakrison, "The Effects of a Visual Fidelity Criterion on the Encoding of Images," *IEEE Trans. Information Theory*, vol. 20, no. 3, pp. 525-536, July 1974.
[22] R. Machiraju, "Analysis, Control and Evaluation of Image Generation in Volume Rendering," PhD thesis, The Ohio State Univ., 1996.
[23] R. Machiraju, A. Gaddipati, and R. Yagel, "Detection and Enhancement of Scale Coherent Structures Using Wavelet Transform Products," *Proc. Technical Conf. Wavelets in Signal and Image Processing, SPIE Ann. Meeting*, vol. 3,169, San Diego, Calif., 27 July-2 Aug. 1997.
[24] S. Mallat, "A Theory for Multiresolution Signal Decomposition: The Wavelet Representation," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 11, no. 7, pp. 674-693, July 1989.
[25] S. Mallat and S. Zhong, "Characterization of Signals from Multiscale Edges," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 14, no. 7, pp. 710-732, July 1992.
[26] S. Muraki, "Volume Data and Wavelet Transform," *IEEE Computer Graphics and Applications*, vol. 13, pp. 50-56, July 1993.
[27] S. Muraki, "Multiscale Volume Representation by a DoG Wavelet," *IEEE Trans. Visualization and Computer Graphics*, vol. 1, no. 2, pp. 109-116, June 1995.
[28] P. Ning and L. Hesselink, "Vector Quantization for Volume Rendering," *Proc. 1992 Workshop Volume Visualization*, pp. 69-74, Oct. 1992.
[29] A. Said and W. Perlman, "A New and Efficient Image Codec Based on Set Partitioning in Hierarchical Trees," *IEEE Trans. Circuits and Systems for Video Technology*, vol. 6, no. 3, pp. 243-250, June 1996.
[30] J.M. Shapiro, "Embedded Image Coding Using Zerotrees of Wavelet Coefficients," *IEEE Trans. Signal Processing*, vol. 41, no. 12, pp. 3,445-3,462, Dec. 1993.
[31] J.-L. Starck, F. Murtagh, and A. Bijaoui, "Multiresolution Support Applied to Image Filtering and Restoration," *Graphical Models and Image Processing*, vol. 57, pp. 420-431, 1995.
[32] H. Tao and R. Moorhead, "Progressive Transmission of Scientific Data Using Biorthogonal Wavelet Transform," *Proc. Visualization '94*, pp. 93-99, Oct. 1994.
[33] P. Teo and D. Heeger, "Perceptual Image Distortion," *Proc. First IEEE Int'l Conf. Image Processing*, vol. II, pp. 982-986, Nov. 1994.
[34] A. Trott, R.J. Moorhead, and J. McGinley, "Wavelets Applied to Loseless Compression and Progressive Transmission of Floating Point Data in 3-D Curvilinear Grids," *Proc. IEEE Visualization '96*, pp. 385-388, Oct. 1996.
[35] J.D. Vilsenaor, B. Belzer, and J. Liao, "Wavelet Filter Evaluation for Image Compression," *IEEE Trans. Image Processing*, vol. 4, no. 8, pp. 1,060-1,053, Aug. 1995.
[36] G. Wallace, "The JPEG Still Picture Compression Standard," *Comm. ACM*, vol. 34, no. 4, pp. 39-41, 1991.
[37] A.B. Watson, Y.Y. Gloria, A.S. Joshua, and J. Vilsenaor, "Visual Thresholds for Wavelet Quantization Error," *Human Vision and Electronic Imaging*, B. Rogowitz and J. Allebach, eds., *Proc. SPIE*, vol. 2,657, pp. 382-392, 1996.
[38] R. Westermann, "A Multiresolution Framework for Volume Rendering," *Proc. 1994 Symp. Volume Visualization*, pp. 51-58, Oct. 1994.
[39] R. Westermann, "Compression Domain Rendering of Time-Resolved Volume Data," *Proc. Visualization '95*, pp. 168-175, Oct. 1995.
[40] R. Westerman and T. Ertl "A Multiscale Approach to Integrated Volume Segmentation and Rendering," *Proc. Eurographics '97*, Budapest, Hungary, Sept. 1997.
[41] Y. Xu, J.B. Weaver, D.M. Healy, and J. Lu, "Wavelet Transform Domain Filters: A Spatially Selective Noise Filtration Technique," *IEEE Trans. Image Processing*, vol. 3, no. 6, pp. 747-758, Nov. 1994.
[42] B.-L. Yeo and B. Li, "Volume Rendering of DCT-Based Compressed 3D Scalar Data," *IEEE Trans. Visualization and Computer Graphics*, vol. 1, no. 1, pp. 29-43, Mar. 1995.
[43] Z. Zhu, R. Machiraju, B. Fry, and R. Moorhead, "Wavelet-Based Multiresolutional Representation of Computational Field Simulation Datasets," *Proc. Visualization '97*, Oct. 1997.

Raghu Machiraju received his PhD from The Ohio State University in August 1996. He is an assistant professor at the MSU/NSF ERC for Computational Field Simulations and an assistant professor of computer science at Mississippi State University. He worked at Control Data Corporation and The Ohio Supercomputer Center for a total of six years from 1986 to 1991. His current research interests include terascale visualization, wavelet methods, and adaptive image synthesis algorithms.

Zhifan Zhu received a PhD in electrical engineering from Mississippi State University in 1995 and an MS in information systems from the University of Essex, England, in 1988. He works in geoscience visualization for the Federal Group of Duke Engineering Services Inc. His research interests include distributed data visualization, feature identification, multidimensional signal processing, and unstructured data visualization.

Bryan Fry received the BS degree in computer engineering from Mississippi State University in 1997. He is currently pursuing the MS degree in computer science at Mississippi State University, where he is a research assistant at the NSF Engineering Research Center. His research interests include wavelets and medical imaging. He was the recipient of 1997-1998 MSU Undergraduate Student Research Award in the College of Engineering.

Robert Moorhead received his PhD in electrical and computer engineering and his MSEE from North Carolina State University in 1985 and 1982, respectively. He received his BSEE summa cum laude and with research honors from Geneva College in 1980. He is the Scientific Visualization Group Leader in the MSU/NSF ERC for Computational Field Simulations and an associate professor of electrical and computer engineering at Mississippi State University. He worked at the IBM T.J. Watson Research Center from 1985 to 1988. He has authored more than 50 papers or book chapters on visualization, image processing, and computer communications. He has received funding from ARPA, ONR, NRL, AFOSR, Army Waterways Experiment Station, the Naval Oceanographic Office, and NASA. Dr. Moorhead was the lead conference chair for IEEE Visualization '97 and is the past conference chair for IEEE Visualization '98.